Opportunities for Advancing Technology-Neutral and Performance-Based Design Methods for the Seismic Design and Regulation of Nuclear Power Plant Structures, Systems and Components

<u>Robert J. Budnitz</u>

Earth Sciences Division
Lawrence Berkeley National Laboratory
University of California
Berkeley CA 94720

September 2010

DISCLAIMER

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231. This project was supported by the U.S. Nuclear Regulatory Commission under a Federal Interagency Agreement with DOE.

TABLE OF CONTENTS

Page **Section**

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

The objectives of this study include a review and summary of the current status of performance-based design and regulatory procedures (the “framework”) to enable nuclear structures, systems and components (SSCs) to withstand earthquakes, and secondly, the identification of potential new advancements and improvements in these performance-based design and regulatory approaches for nuclear energy systems that can enhance design and regulation, stimulate innovation, and provide cost efficiencies, among other benefits. Specifically, this report discusses how the current “framework” uses probabilistic-based performance targets, and contrasts them with the well-established traditional approaches.

ACKNOWLEDGEMENTS

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231. This project, known at NRC as project #N6509, was supported by the U.S. Nuclear Regulatory Commission under a Federal Interagency Agreement with DOE.

It is a pleasure to acknowledge the strong support of Dr. Annie Kammerer of the NRC Office of Nuclear Regulatory Research, whose insights into NRC's regulatory needs has provided the fundamental rationale for this work. Others at the NRC, in particular Dr. Nilesh C. Chokshi, also provided valuable technical insights. Major parts of the overall NRC project of which this work has been a part are being carried out by Prof. Bozidar Stojadinovic of LBNL and UC-Berkeley, and Prof. Boris Jeremic of LBNL and UC-Davis. Valuable discussion with both of them provided important benchmarks for how to approach this problem. The project began under the direction of Dr. David B. McCallen, then at LBNL and now at the Lawrence Livermore National Laboratory; his initial draft of parts of Section 5 of this report provided a jumping off point for what later became a very different document, and his crucial insights into how structural-engineering codes actually work in practice was of great help.

A review committee chaired by Dr. Robert P. Kennedy provided valuable help in the initial stages; others on that initial review committee were Donald P. Moore of Southern Nuclear Operating Company, Prof. Gregory Fenves of the Univ. of Texas-Austin, Dr. Chokshi, and Prof. Jeremic. A second meeting of the review committee provided additional important input; it consisted of Dr. Kennedy (chair), Mr. Moore, Dr. Chokshi, Dr. James J. Johnson (consultant), and Dr. Antonio R. Godoy (International Atomic Energy Agency, retired.)

ACRONYMS AND INITIALISMS

ANS	American Nuclear Society
ASCE	American Society of Civil Engineers
CDF	Core Damage Frequency
DBE	Design Basis Earthquake
DF	Design Factor
DOE	US Department of Energy
DRS	Design Response Spectrum
EPRI	Electric Power Research Institute
FOSID	Frequency of Onset of Significant Inelastic Deformation
LS	Limit State
LWR	Light Water Reactor
NRC	US Nuclear Regulatory Commission
PRA	Probabilistic Risk Assessment
PSHA	Probabilistic Seismic Hazard Analysis
SDC	Seismic Design Category
SF	Safety Factor
SPRA	Seismic Probabilistic Risk Assessment
SSC	Structure, System, or Component
SSHAC	Senior Seismic Hazard Analysis Committee
UHRS	Uniform Hazard Response Spectrum

EXECUTIVE SUMMARY

The information contained in this report was prepared to document one portion of a research project conducted for the U.S. Nuclear Regulatory Commission's Office of Nuclear Regulatory Research. Objectives of this study include first a review and summary of the current status of performance-based design and regulatory procedures to enable nuclear structures, systems and components (SSCs) to withstand earthquakes, and secondly, identification of possible new advancements and improvements in these performance-based design and regulatory approaches for nuclear energy systems.

Specifically, this report describes the current status of design and regulatory approaches (the "framework") in use in the U.S. for nuclear power plant SSCs to enable them to withstand earthquakes, discusses how these current approaches use probabilistic-based performance targets, contrasts them with the well-established traditional approaches, and identifies potential performance-based advances beyond today's status that can enhance design and regulation, stimulate innovation, and provide cost efficiencies, among other benefits.

The report begins with a broad conceptual introduction to the traditional approach to design of safety-critical SSCs in nuclear power plants to withstand earthquakes, including a discussion of the major concepts. The report then outlines the principal characteristics of probabilistic analysis, performance criteria, and performance-based design.

This leads into a discussion of the several different characteristics of a "framework" for the form and content of the performance criteria, the acceptance criteria (including the confidence level), the analytical methods, and the design criteria. For each of these, the developer of the "framework" has several options, which are explored and analyzed to provide background that can help a decision-maker decide which options to select. The options range from some that represent only modest departures from the traditional engineering design approach to others that represent quite radical departures.

Next, an overview is provided of today's "framework", as embedded in the consensus American Society of Civil Engineers code ASCE 43-05 and in part in the US Nuclear Regulatory Commission staff's guidance. This ASCE 43-05 framework uses a target performance goal as its major attribute, and provides detailed design and analysis guidance that, if executed properly, aims at providing confidence that an individual SSC will meet that performance goal. How this works in practice is discussed in some detail, although the specific technical details of the engineering guidance are not repeated here. The performance implications of following the framework are also discussed. Specifically, if the framework is successfully implemented, there should be confidence that an SSC designed using it will have less than about a 1% probability of unacceptable performance for the Design Basis Earthquake Ground Motion, and less than about a 10% probability of unacceptable performance for an earthquake with 150% of the Design Basis Earthquake Ground Motion.

The report then discusses how today's ASCE 43-05 framework lines up with respect to the various options discussed earlier in the report. This leads to some observations about where certain limitations exist in that framework. This in turn leads to the final section of the report, where several possible paths forward beyond today's framework are discussed. These range from some narrow but important proposals (such as designing SSCs by accounting properly for their safety functions within specific PRA accident sequences, rather than requiring each SSC individually to meet the same performance target, and such as using a more realistic definition of earthquake-caused "failure"), to the broad and controversial proposal to use a fully risk-based (that is, PRA-based) criterion for deciding how to design a given SSC to resist earthquake loads, or alternatively to permit a fully performance-based approach in which the designer need follow no specific design codes or rules, but must perform a very extensive analysis to demonstrate that a pre-set regulatory performance target is met with adequate margin.

It is recognized that this latter idea is probably not feasible at this time, because the state of the art of neither design nor analysis is adequate to provide a sufficient demonstration that adequate safety is achieved. However, proposals are presented for steps forward along the path toward this ideal objective, accompanied by a discussion of why some of the most fully risk-based ideas are unlikely to be adopted any time soon.

1.0 Overview

The information contained in this report was prepared to document one portion of a research project conducted for the U.S. Nuclear Regulatory Commission's Office of Nuclear Regulatory Research. Objectives of this study include a review and summary of the current status of performance-based design and regulatory procedures to enable nuclear structures, systems and components (SSCs) to withstand earthquakes, and secondly, identification of possible new advancements and improvements in these performance-based design and regulatory approaches for nuclear energy systems.

Specifically, this report describes the current status of design and regulatory approaches (the "framework") in use in the U.S. for nuclear power plant SSCs to enable them to withstand earthquakes, discusses how these current approaches use probabilistic-based performance targets, contrasts them with the well-established traditional approaches, and identifies potential performance-based advances beyond today's status that can enhance design and regulation, stimulate innovation, and provide cost efficiencies, among other benefits. It begins with a broad conceptual introduction to the traditional approach including a discussion of the major concepts.

2.0 The traditional engineering design approach to a "successful" design

For structural systems generally, the time-honored approach to design embodies the simple principle that every engineering student learns at the beginning, which is that after having defined the "load" that the design must carry, achieving a successful design requires that *the Capacity C must exceed the Demand D:*

$$C > D$$

When this is true, the design is "successful," meaning that the engineered system can withstand the load and *continue to perform whatever function it is designed for.*

Of course, every student also learns that this can only be realized in practice *if the load is well-defined*, including loads in off-normal conditions. Crucially, to ascertain whether in fact C > D, the engineer must be able to *analyze both C and D with sufficient accuracy and adequate confidence,* or *must gain the needed confidence using a different approach that does not rely on detailed analysis.*

To call a design "successful", it needs to meet three criteria: (a) C must indeed exceed D in fact; (b) the designer must have adequate confidence that in fact C exceeds D; and (c) the approval of the appropriate regulatory or licensing authority must be obtained. These three criteria interact, and meeting (a) is not sufficient. We will discuss the three criteria together in what follows.

The above short discussion is, of course, too succinct to be of any practical use, because even if the load is well-defined, it begs a number of crucial questions, the most important of which are:

- What is the actual demand?
- What is the actual capacity?
- How is the phrase "to perform its function" defined in practical terms?
- What are the decision criterion and the decision process in practice for declaring a design "successful" and hence for allowing the project to proceed?
- There is always uncertainty. How is it handled in the decision process?

For designs that are accomplished today, although it is always necessary to define the function, analyzing both C and D is not always necessary. This is because of the almost universal use of consensus codes and standards. Another early lesson that every student learns is that if an engineer faithfully follows the established design code(s) for a given design problem, the design will be "successful", provided that the <u>scope</u> of the established design code appropriately matches the <u>scope</u> of the design under consideration. That is, given that the desired function is well defined, and that the load is defined for the desired function, standard practice implies that <u>following the code will assure a successful design</u>. The underlying concept is that the consensus code committee's deliberations, as embedded in the code itself, have taken care of all of the issues that the individual engineer would need to resolve in the absence of the code.

This conceptual approach is in fact the basis for how most engineered structures and systems are designed today. The broad societal objective is to assure that all designs are "adequate" while recognizing that individual engineers at the bench are not always well enough trained or competent enough in executing their work to account for everything that needs to be considered. Hence "following the code" is embedded in the culture of practical engineering today, and appropriately so.

This being said, there are two universal problems: first, *the function needs to be well defined* and second, *the "load" (that is, the "demand D") must also be determined,* either accurately or by following a prescribed methodology, after which "following the code" will supposedly assure that C > D.

2.1 Demand

An ever-present complication is that the Demand is usually not easy to specify, except in cases in which the codes themselves provide a prescription for defining it, or in some straightforward situations such as when the demand is a simple constant static load for a simple system. However, for a complex load such as the category of earthquake loads, one cannot usually conclude that C > D for all earthquakes that might arise: the Demand will vary from one earthquake to the next, and some possible earthquakes are much larger than a typical design can withstand. Therefore, in fact the problem of specifying the Demand comes down in part to asking whether the successful outcome must occur for earthquakes of a specified "size" corresponding to a given "return period" or "annual frequency", or perhaps for earthquakes of a larger "size" corresponding to a longer "return period" or a smaller "annual frequency."

This complication is not unique to earthquake loads. Other natural hazards (for example, high winds) and some man-made hazards also exhibit this complication.

Because it is understood that there is a risk of "failure" of the system for earthquakes very much larger than those designed for under the code, this issue becomes one of risk management: *How small must the residual risk be to be acceptable?* For loads like earthquake loads, this is an issue usually taken on directly by the consensus code committees (or the regulators), by specifying how the Demand to be used in the design is to be derived, and in what form. Typically, the specification is in terms of how large the design earthquake must be, or in terms of some parameterization related to that concept. This specification intrinsically contains within it a decision on how much residual risk is acceptable, in terms of the frequency of earthquakes with a "size" exceeding whatever performance cutoff has been chosen. Sometimes this risk decision is explicitly stated and explained by the code or regulation, but not always.

2.2 Capacity

The analysis of actual Capacity can also be difficult if a realistic result is needed. Except for simple systems with a well-characterized failure mode, there will inevitably be a lot of uncertainty in any such analysis. Accounting for the additional observation that there can be a good deal of variability in fabricated items that are seemingly identical (due to materials variations, variability in fabrication and maintenance, etc.), this means that highly accurate analyses of Capacity are usually beyond the capability of all but the most skilled engineers. Often, the analyst's knowledge of the Capacity of a structure or component is presented as a lower bound, such as, "There is high confidence that the Capacity is at least as large as X." Sometimes the "high confidence" can be expressed numerically (for example, 85% or 95% confidence). Less commonly, the analyst's knowledge of the Capacity is expressed as a *fragility curve* that gives the probability of "failure" as a function of the "size of the load" or of the "demand" (an example is in Figure 1). However, to derive a fragility curve requires simplifying the demand in the form in a one-dimensional parameter (the abscissa in Figure 1), which itself is almost always a simplification.

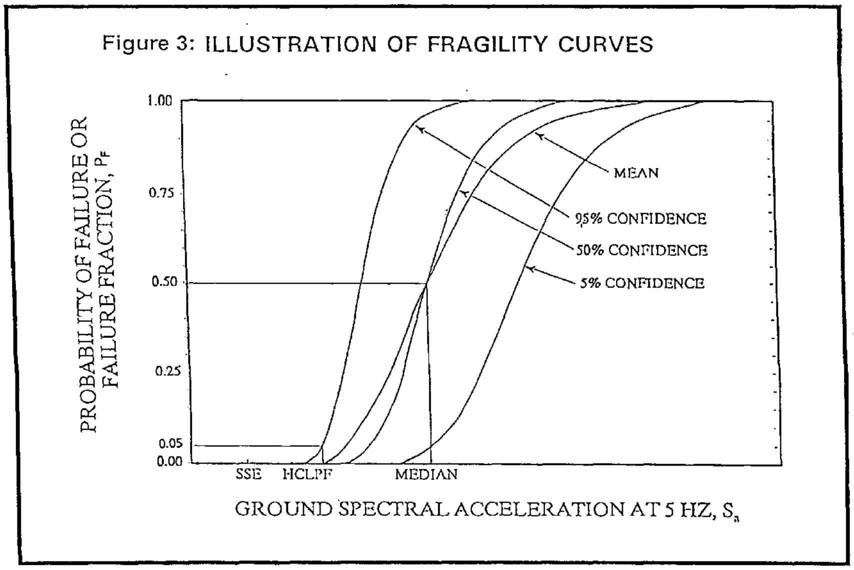

FIGURE 1
Illustration of typical fragility curves for a component

2.3 Functional performance

In contemplating what it means for C to exceed D, another vital issue is what is meant by the phrase "continue to perform whatever function it is designed for." For a simple civil structure like a single-family residence, current earthquake codes seek to assure protection of life and limb, whereas for a vital community facility like a hospital the earthquake codes seek to assure that it can continue to function after the earthquake. And for a nuclear power plant, the NRC regulations and the consensus codes seek to assure that each vital SSC can continue to perform its safety function after the earthquake, so as to protect against release of radioactivity from the reactor plant to the environment. In the nuclear reactor regulatory environment, a "failure" criterion is typically selected that is quite conservative relative to the actual failure to perform the required safety function.

2.4 Uncertainty, conservatism, and margins

Another complication is that of uncertainty. If assuring that C > D requires the ability to *"analyze both C and D with sufficient accuracy and adequate confidence"*, then the analytical basis for determining both C and D comes into play, in the face of both uncertainty and variability.

There are many different issues involved in understanding the uncertainties in achieving the desired outcome. Some of the uncertainties are in the data supporting the analyses, some are in the inherent variabilities in the loads, some are in the material properties and construction methods, and some are in the models used for the analysis, which are always an abstraction of reality rather than fully realistic. The consensus code committees and government regulatory agencies must deal with these uncertainties. They do this by introducing conservatisms in the analyses supporting the requirements in the codes and in the data used in these analyses, along with conservatisms in the analyses required of the design engineer. Thereby, conservatisms are inserted into the engineer's design process.

The conservatisms can take many different forms. While these are often thought of a "safety factors", they usually are not exclusively in the form of a simple multiplicative Safety Factor (SF) incorporated into the D/C comparison, namely a factor that either reduces the apparent Capacity:

$$C/SF > D \text{ (here } SF > 1).$$

or increases the apparent Demand:

$$C > D * SF \text{ (here } SF > 1).$$

Besides the straightforward use of Safety Factors, several other approaches are used in the traditional design codes to achieve an adequate margin to account for uncertainty. Some of them include forcing the analyst to assume pessimistic values for certain of the parameters in the analysis, or to use simple models known to be conservative, or to ignore certain mitigating phenomena, or to specify one or more clearly biased assumptions that will result is an implicit or explicit SF, in addition to forcing one to design for a greater load, or to assume a smaller capacity than will realistically occur.

All of the above discussion is quite general. Because we wish to concentrate on design against earthquakes, from here on out the discussion will explicitly concentrate on that class of design problems.

2.5 Deterministic methods and assurance achieved

Generally, all of the above are embedded in the design codes and in regulations through specified methods, so-called "deterministic" methods, used for design and for analysis.

If everything works as it is supposed to work, then there is assurance that C > D. This can take two different forms. Usually, the approach will provide high assurance that the C will exceed the specified D. Alternatively, sometimes the goal is to provide assurance that the Capacity will exceed a Demand somewhat larger than the specified D, but with somewhat lower assurance. Or both.

If the latter, which in fact is not typically part of the traditional process, two crucial questions can be posed: *How much larger? And with how much "(lower) assurance"?* The traditional design approaches used in nuclear safety design against earthquakes in the past usually did not answer these questions, and in fact they were typically not able to, because not enough information was developed to enable these questions to be answered.

Finally, but not least important, the traditional design approaches embed or require the use of prescriptive methods, which specify how the design is to proceed and how the analysis is to be performed. Such prescriptions discourage or stifle innovation, and purposely so in the name of assuring a measure of uniformity across the wide spectrum of designers and analysts who practice the trade at any given time.

Furthermore, in the hands of the routine practitioner the traditional approach often leaves the underlying assumptions and limitations obscure or unexpressed. The routine practitioner often has no idea about what the underlying assumptions and limitations might be.

2.6 Other limitations

Two other widely recognized limitations should also be mentioned. The first is that, for most engineered systems, the traditional approach to design to withstand earthquakes in nuclear plants does not involve analyses that work out the *realistic performance* of the system under various *off-normal scenarios*, even for scenarios that are intended to be encompassed by the design, and all the more so for scenarios more severe than that. The design codes and regulations usually do not require such realistic analysis, the system's owners typically do not ask for it, and in fact the methods for working out such an analysis often are not in common use.

The second limitation is that generally there is no way to know the *fragility curve*, namely what is the *probability of reaching the undesired endpoint* (that is, true "failure") as a function of the "size" of the load when the load exceeds the design demand load, for which presumably there is very high assurance of successful performance. Again, such an analysis is usually not required by the traditional codes and regulations, is not usually requested by the owner, and sometimes cannot be done anyway using commonly available methods or data. But, as should be obvious, such information is of great importance, if not for each SSC then surely for the ensemble of all such that together comprise the safety equipment and structures of a nuclear power plant.

Despite the several limitations described briefly above, the traditional approach clearly "works" in general --- a very large number of structures and components are deployed successfully in the hundreds of operating nuclear power plants around the world. However, the several limitations touched on above seem to cry out for a more advanced approach that can overcome some or all of them.

And in response, in recent years the modern philosophy best described as *design to achieve an explicit performance target* has come to maturity, coupled with the advent of *probabilistic analysis methods* as the vehicle for overcoming many of these limitations. This modern philosophy also interacts with another modern trend, the trend toward *"performance based design"*, which of course requires explicit performance criteria on which to base the design and its analysis.

The next Section discusses these concepts separately, and then explores the nexus among them.

3.0 Probabilistic analysis, performance criteria, and performance-based design in nuclear power plant seismic design and regulation*

3.1 Introduction – performance-based design, advantages and challenges

The major advantage of performance-based design is that in principle it can allow a designer great flexibility in finding innovative ways to achieve the desired performance. If the desired performance is simply to achieve very high confidence that C > D for a specified earthquake Demand, the philosophy is relatively simple to implement, at least in principle. If, however, as discussed below, the desired performance includes specified assurance of performance for earthquake loads where "failure" is at least admitted as encompassed within the design target (albeit at defined low probabilities), then implementing the modern philosophy described here is by no means simple. The discussion that follows will explain why.

3.2 Introduction – probabilistic analysis

Especially for earthquake loads beyond where very high assurance of adequate performance is achieved, the fundamental basis for the modern trends to be covered in this section is the observation that many of the most important questions that facility owners, government regulators, and the general public want to understand (and have a right to understand) about the safety of large facilities are *questions that can only be framed and analyzed probabilistically,* meaning can only be analyzed using probabilistic methods.

As mentioned, the key issue involves performance for loads well beyond the design loads, and in practice this usually means *for loads greater than those for which very high assurance of adequate performance can be achieved*. The occurrence of these loads is not an every-day event (such as always being present, or for example being present with certainty every time the system starts up or shuts down), or they would be within the design. Rather, their occurrence

* The discussion up to here has not been focused entirely on earthquake-safety issues for power reactors. Starting in this section, the discussion will concentrate on the issues of design, analysis, and regulation of nuclear reactors and their SSCs to withstand earthquake loads.

has a probabilistic character, in terms of their probability of occurrence and also in terms of their “size”, however size is defined. *Also, the response of the system has a probabilistic character, in that the probability of off-normal undesired response is almost never a zero-one effect --- instead it almost always involves a probability of “failure” as a function of the “size” of the above-design load.*

Hence answering questions about the response of a system beyond its design, in the region where “failure” occurs, intrinsically must involve probabilistic analysis using probabilistic descriptions of the phenomena and probabilistic data.

A second key point is that public policy typically recognizes that there is no such thing as “perfect safety” – it is necessary to accept that for possible off-normal events of very large severity but very small likelihood, such as very large earthquakes, there will be a risk of “failure” and hence a safety risk, to the public, to the facility, and/or to the environment. That is, we cannot achieve perfect protection, albeit we can strive for it. In this light, *to inform public policy on how much protection to seek*, taking account of the costs and other issues involved with achieving a given level of protection, *involves asking and answering probabilistic questions about the risk*. These questions boil down to three, known in the literature as the “risk triplet” of questions: “What can go wrong?”, “How likely is it?”, and “What are the consequences?” -- all as a function of the “size” of the postulated severe event.

3.3 Scenario analysis

Notice that the answers to these questions must necessarily be provided in the context of an ensemble of off-normal scenarios, because the answer to the first question, “What can go wrong?”, is a set of distinct scenarios each of which requires analysis separately. *Thus the rest of the analysis must be scenario-based.* One matter for public policy then becomes deciding which scenarios are to be included in the safety envelope requiring design solutions, which scenarios are to be studied and understood (and perhaps planned for) even though they are nominally beyond the design envelope, and which scenarios are considered so unlikely that they need not be accounted for. This has implications both for the types of scenarios (“Are only certain types of hazards or other off-normal events to be considered in the design?”) and for the probabilities of them (“Can scenarios less probable than a frequency of once is a zillion years be excluded?”, and if so, does once-in-a-zillion mean every 100 years or every 10,000 years or every 1,000,000 years?)

3.4 Probabilistic performance objectives, probabilistic acceptance criteria, and probabilistic analysis methods

Today, it is both feasible and increasingly common for consensus code committees and regulatory agencies to establish performance objectives for an engineered system to withstand earthquakes that are framed in probabilistic terms. This is made possible by the advent and indeed the widespread use of probabilistic analysis methods and even some probabilistic design methods. For an external hazard like earthquakes, these probabilistic analysis methods address different aspects of the overall problem, as follows:

o *Probabilistic hazard analysis methods* are used to describe probabilistically the likelihood of hazards of different “sizes” as a function of the size and of other relevant parameters. The archetype of this is *probabilistic seismic hazard analysis (PSHA).*

o *Probabilistic fragility analysis methods* are used to work out the probability of “failure” as a function of the “size” of the load and the cause of the failure. Here it is vital to define in a precise way both what is meant by “failure” and how the load is to be described. The archetype of this is *probabilistic seismic fragility analysis*.

o *Probabilistic system response analysis methods* work out how the failures of individual elements or parts of a complex system come together to produce a failure of the larger system. Again, the description of this overall failure is described probabilistically in terms of a combination of the lower-level failures. The archetype of this is *probabilistic risk assessment (PRA) systems analysis for a nuclear power reactor*, typically using event-tree and fault-tree methods.

In addition to these 3 methodologies that deal with parts of the overall analysis problem, the methodology of PRA integrates all 3 of the above together into an analysis that provides results in the form of risk, probabilistically expressed with uncertainties. For our purposes here, the archetype is *seismic PRA (SPRA)*, which has been used to study dozens of nuclear power plants and other facilities. This is not the place for a lengthy discourse on SPRA, of course. However, a major attribute of PRA is uncertainty analysis, which is very much related to the analyst’s confidence in the results.

Given the availability and the increasingly widespread use of the above types of probabilistic methods, it is feasible (as mentioned above) for consensus code committees and regulatory agencies to establish *probabilistic performance criteria* for a given system. And they are increasingly doing so for scenarios involving earthquakes as well as scenarios involving other causes of upset conditions. Furthermore, it is now becoming feasible to specify a range of performance criteria for different types of undesired outcomes. An example is seeking at least less than a frequency F_1 (say, 10^{-k} per year) that the system will sustain even minimal damage, less than a smaller frequency F_2 (say, 10^{-m}/year) that damage will not exceed certain undesired thresholds such as a repair cost or downtime, and less than a still smaller frequency F_3 (say, 10^{-n}/year) that severe damage close to destruction will occur. Here, as indicated, $F_3 < F_2 < F_1$. A crucial aspect of such a specification must be specifying how much confidence (median confidence? 95% confidence? 99% confidence?) the analyst needs in the conclusion(s).

3.5 The risk profile

The advantage of the above, if executed well, is that the analyst, the owner, the regulatory decision-maker, the political policy maker, and the general public can all ask and obtain answers to a range of questions that were inaccessible before. *The most important of these are questions whose overall umbrella concept is the risk profile.* The risk profile describes the probabilities and consequences of all of the important off-normal scenarios that might occur at a facility, including a description of the analyst's uncertainty in his state of knowledge about both the probabilities and the consequences.

And crucially, the advent of risk-analysis methods enables the public policy maker to establish regulations that can realistically expect to require that certain risk-profile outcomes be achieved, probabilistic though they must be.

3.6 Realistic analysis, the essential backbone – the essential ingredient, the *sine-qua-non*

All of the above, under the rubric of probabilistic analysis and performance objectives, as desirable as it may seem to be in terms of overcoming some of the major limitations of the traditional engineering approaches to design and regulation, cannot be achieved fully without the ability to perform *realistic analysis* of the complex system under review. Notice also that the realistic analysis must be done for the whole spectrum of off-normal scenarios, one scenario at a time, because that is the only way that the risk can be described for any such system. Without the use of realistic scenario-based analysis, a realistic description of the risk profile is unavailable, as is the comparison of that risk profile with the performance objective(s) to determine whether the system is "safe enough." [*]

In practice, less than fully realistic analysis can be prescribed as acceptable. There is a compromise here. The issue is one of deciding where to introduce the conservatism that a decision maker needs: in the design requirements, in the requirements for analysis, or in the decision framework once the design and the analysis have been completed and a go/no-go decision is required.

In practice, for a probabilistic approach to be successful, all of the above need to be incorporated into a consistent "framework" that is tied together. Several options exist for such a framework, which will be discussed in the next section.

[*] Of course, analysis that is less than realistic, if demonstrably conservative in character, can always be used to demonstrate a bound on the risks associated with various scenarios. However, the trend is increasingly toward making the analyses more and more realistic, in order to derive the maximum benefits from the approach.

4.0 The major options for the "framework" [Options for the form and content of the performance criteria, acceptance criteria (confidence level), analytical methods, and design criteria]

With the discussion in section 3.0 as background, several options present themselves, a combination of which could serve as <u>performance criteria</u>, <u>acceptance criteria</u>, <u>prescriptions for the analytical methods</u>, and <u>design criteria</u>. Taken as a unified set, these would then serve as the basis for structural/seismic design, analysis, and regulation in the nuclear-power-plant arena. *We will call this unified set the "framework"*.

In this section, the various options available for the framework will be laid out and their attributes described.

It is important to note that a regulator or a consensus code committee needs to select among these options taking the whole picture into account – that is, it is crucial to account for the ramifications of all of these options as a <u>set of decisions</u> is being made.

<u>4.1 Options for the performance criterion and the level of confidence</u>
As a matter of logic, the entire probabilistic framework must rest on a criterion for acceptable vs. unacceptable performance that is expressed in probabilistic terms. As explained above in section 3.0, there is no other way to begin.

<u>A) The Criterion</u>

There are two Options for the form of the probabilistic performance criterion. It must be defined in terms of either

<u>(Option a)</u> an *annual frequency of unacceptable performance*, meaning a risk-type criterion, namely an integrated risk that captures an integration over the probabilistic hazard and the probabilistic fragility curve to give a criterion in the form of an *overall frequency of "failure"*. An example could be that the overall likelihood of seismic-caused unacceptable performance ("failure"), expressed as a frequency, must be less than 10^{-5} per year.

or

<u>(Option b)</u> some *surrogate* for the above. An example could be that the probability of unacceptable performance must be less than X% for a prescribed earthquake that would occur with frequency F per year (e.g., a 1% probability of "failure" for a 10^{-4}/year earthquake).

The selection of the above criterion is a policy choice, not a technical choice, involving how much protection is being sought.

B) The Confidence Level

We assume that the *required degree of confidence* is also prescribed, usually in terms of a level of confidence that the criterion is met (median confidence? 95% confidence? etc.) This is another policy choice, not a technical choice. Crucially, this choice is closely related to the uncertainty in the analysis, which itself is closely related to how much "margin" a decision-maker (a consensus code committee, a regulator, or the owner of the facility) requires in order to conclude that the design will be "safe enough."

The use of this approach necessarily means either that the analyst must be required to perform a robust uncertainty analysis, or that a *surrogate* is provided (by the code committee or the regulator) and used to capture the same notion. An example of a surrogate could be that a median frequency for the risk is required to be demonstrated by analysis, and then some prescribed extra margin is required in order to assure that the acceptance criterion is met --- say, an extra factor of 2 or 4. The basis for the numerical value of this extra margin presumably arises from knowledge (developed and used by the code committee or the regulator) from actual experience or analysis. This experience or analysis must, of course, be applicable to the case at issue, which places a burden on the code or regulation to be explicit about the scope – what's "covered" and what's not -- so as to avoid the danger that the code or regulation is applied where it is not applicable.

4.2 Options for the form of the performance criteria and of the analytical methods

Given that the Criterion and the Confidence Level have been established, there are a number of options for the *form* of the rest of the "Framework." We will outline them next. It is vital to note again, however, that a decision-maker who selects among the 3 options below needs to consider the entire set of decisions on the options as a package, because they all interact.

Notice that because each option below has two or three branches, there is a 2 x 2 x 3 set of possibilities here, or 12 possibilities.

Option One: Performance -- "Success" vs. "Failure" (Realistic or Conservative)

Is the "performance" in the probabilistic performance criterion a realistic definition of "success" vs. "failure", or is a conservative end-point chosen? (An example of the latter is to select the onset of modest inelastic deformation as the conservative definition of seismic-induced "failure" for a shear wall.)

Option Two: End-Point: One SSC at a time, or Accounting for Systems Analysis

For a nuclear reactor, given that the actual end-point of concern in terms of safety assurance is generally taken as preventing or mitigating “core damage”, is the end-point of the design criterion under consideration the design of a single SSC (one at a time), or is use made (or allowed) of a systems analysis structure in which it is the frequency of “failure” of a full accident scenario that is to be kept below the target?   At least three different options exist here: one SSC at a time? or one such scenario at a time? or by an integrated accounting for the full ensemble of accident scenarios?

Option Three: Analysis – Realistic or Requiring Conservatisms

Is the analysis to demonstrate compliance required to be “realistic”, or is the analysis method specified in terms of certain conservative analytical methods or the use of specified conservative data?

As mentioned above, when these 3 options are considered together, the total number is 2 x 2 x 3 or 12 possibilities for these 3 options.

4.3 Options for the design side

The discussion above covers the broad options available to the standard-setter, in a probabilistic “framework”, in terms of the performance criteria and how they are to be met. But that aspect of the framework is only part of the whole picture. Next one must consider the requirements or guidance to the designer.

Here, the options are comparably broad, ranging from true performance-based design to design using entirely prescriptive design rules.

Option A: True Performance Based Design

In this option, the designer need not follow any prescriptive design rules, but is free to execute the design however he/she wants to or needs to, with the constraint that of course the design must meet all of the requirements set down for demonstrating compliance.

This option places the heaviest burden not on the design but on the compliance analysis. Indeed, accounting for uncertainties and establishing the requisite degree of confidence is a major task, fraught with analysis difficulties: both difficulties with demonstrating the completeness of the scope of the analysis, and documentation difficulties. This burden is so great that if this approach is used, a bit of wise advice to the designer is that care must be taken with analyzability – indeed, the old dictum comes into play that “analyzability must be a design criterion”, because a design that cannot be analyzed well enough will not be approved (should not be approved!) and therefore cannot be built.

A major advantage of Option A is that it imposes few constraints to stifle innovation in design, except the constraint that the design must be analyzable.

Option B: Specified Deterministic Design Rules with a Performance Criterion

In this option, the designer is constrained by specified design rules, embedded either in a consensus industry code or a government regulation. However, the design must be analyzed to show compliance with a performance criterion along the lines discussed above. The rationale for this Option is that designers need the constraints and guidance that the consensus code rules provide, because the practicing designer at the bench should not simply be "turned loose" in design space without detailed rules.

Some of the advantages of Option A's flexibility are lost with Option B, because some innovative design ideas simply cannot be accommodated by the established design rules in the consensus codes. By their nature, the codes must lag behind the latest innovations. These disadvantages are compensated by the advantage that use of the design rules provides a constrained "envelope" within which the design is executed – and (presumably) within which the design will then perform.

As with Option A, there is the constraint that the design must be analyzable to show compliance with the performance criterion, although some variants on this Option could provide specified analysis rules too.

Option C: Specified Deterministic Design Rules without an Explicit Performance Criterion

In this option, as with Option B, the designer is constrained by specified design rules, embedded either in a consensus industry code or a government regulation. However, the designer need not perform an analysis to demonstrate compliance with a performance criterion. Instead, *compliance is achieved (and implicitly demonstrated) simply by having followed the design rules*.

The development of the design rules, in turn, can be characterized by one of two different approaches – different "logics", as it were:

(Logic a) In this approach, the regulator or consensus code committee has developed an explicit performance target that the designs must meet, has published and explained what the target is, and has developed the design rules so that the target "comes true" if the rules are followed.

(Logic b) In the other approach, nothing as explicit as the above is written down. Indeed it may or may not be explicit in the minds of the regulator or the code committee when the rules are developed. This is, of course, closer to the more traditional “deterministic” approach to design and regulation that has characterized the profession from time immemorial.

The choice between these logics, or choosing something in-between, obviously depends in part on whether the decision-maker (the code committee or the regulator) believes that it can develop explicit “deterministic” design rules that, if met, will assure that the performance target is realized. If so, then “logic a” can be selected and defended. Otherwise, with no way to defend a claim that an explicit performance target is likely to be met for a particular design, one is driven to “logic b.” However, since in-the-end every approach seeks to meet a performance target, even if it is implicit target, the “logic b” approach must incur a penalty. That penalty is that inevitably additional conservatism must be embedded in the design, either through a conservative implicit performance target, conservative design rules, or conservative acceptance criteria.

Using this Option C clearly places a heavier burden on the consensus code committee or the regulator, because their decision on the design rules is effectively all-powerful --- a design is not checked by analysis against the performance criterion that presumably underlies the design rules, and this is true whether the criterion is explicit or implicit.

4.4 Summary discussion on the options

As mentioned above, it is important to note again that a regulator or a consensus code committee needs to select among all of these Options taking the whole picture into account – that is, it is crucial to account for the ramifications of all of these Options as a set of decisions is being made.

In practice, a consensus code committee or a regulator is free to adapt any of these options, or to mix-and-match. As we shall see, the actual codes and regulations that are in place today can best be characterized in that way.

A summary overview of the several different options is presented in Tables 1, 2, and 3, sorted into the three different areas: (i) options for the probabilistic performance criteria, (ii) options for the form of the performance criteria and for analysis of performance, and (iii) options for the design requirements.

Although there has been an evolution over time, only one combination of these options is endorsed today by the NRC to guide seismic design for nuclear power plants and to govern how the analyses must be performed.

In the next Section (Section 5), we will describe the approach used in the existing design standards, analysis standards, and regulatory standards for U.S. nuclear power plants, and we will see how they fit together as a “framework.” Then, in the follow-on Section 6, we will identify which of the several options for such a framework is represented by today’s approach.

TABLE 1
Options for the probabilistic performance criterion and for the confidence level

probabilistic performance criterion	confidence level
annual frequency of unacceptable performance	explicit high confidence
	median confidence + extra margin arising from somewhere else
a surrogate	explicit high confidence
	median confidence + extra margin arising from somewhere else

TABLE 2
Options for performance and for analysis

Options: Performance	Options: End-Point	Options: Analysis
Realistic Success v. Failure	full ensemble of accident scenarios	realistic analysis
		conservative analysis
	one accident scenario	realistic analysis
		conservative analysis
	a single SSC	realistic analysis
		conservative analysis
Conservative Success v. Failure	full ensemble of accident scenarios	realistic analysis
		conservative analysis
	one accident scenario	realistic analysis
		conservative analysis
	a single SSC	realistic analysis
		conservative analysis

TABLE 3
Design side options

True Performance-Based Design	Specified Deterministic Design Rules with an Explicit Performance Criterion	Specified Deterministic Design Rules without an Explicit Performance Criterion

5.0 Overview of the approach used in existing earthquake design and analysis codes and standards for nuclear power plants

5.1 Introduction

Today, the basic procedures for seismic design of nuclear facility structures, systems and components (SSC) are provided in the American Society of Civil Engineers Standard ASCE / SEI 43-05 (Ref. 1), an important professional consensus standard. The intent of the ASCE 43-05 Standard is to produce engineered designs that achieve an acceptable *target seismic risk goal*, defined in terms of an *annual probability of seismically induced unacceptable performance*. This is accomplished by meeting established annual probability of unacceptable performance *target performance goals*, while simultaneously not exceeding *specified limit states*. The way a design is accomplished differs depending on the type of facility, and facilities are categorized into different *seismic design categories.*

There are 5 different terms here representing different concepts, each of which needs to be understood:

- target seismic risk goals;
- annual probability of seismically induced unacceptable performance;
- target performance goals;
- limit states;
- seismic design categories.

ASCE 43-05 covers the design of a wide range of nuclear facilities besides nuclear power plants, using a graded approach, and the discussion that follows will incorporate some features of how the standard deals with these other facilities. However, our emphasis here will be on nuclear power plants.

5.2 Seismic Design Categories, Limit States, and Target Performance Goals

Seismic Design Categories and Limit States: The graded approach in ASCE 43-05 introduces different levels of conservatism depending on the facility; this approach is intended to ensure that the levels of conservatism are consistent with the functionality and human hazards associated with a particular facility. To that end, every nuclear facility is assigned one of five "Seismic Design Categories" (SDCs), using guidance and following the SDC descriptors found in the American Nuclear Society Standard 2.26 (Ref. 2), and shown in Table 4. In addition, a functionality descriptor, called a "Limit State", is assigned to each facility being designed, again using guidance in ANS 2.26, and shown in Table 5. The choice of the SDC and LS for a given facility would normally be the prerogative of the facility's owner, but for vital SSCs in nuclear power plants a particular SDC and LS assignment has been made by the US NRC.

ASCE 43-05 only covers SDC categories 3, 4, and 5. SDC categories 1 and 2 are addressed in a different standard, the building standard contained in ASCE 7 (Ref. 3) for typical general building structures, but because facilities of this type are outside our scope here, these SDCs will not be discussed in detail further here.

The 5 Seismic Design Categories and 4 Limit States described in ASCE 43-05 are listed in Table 6, along with the appropriate standard (ASCE 43-05 or ASCE 7) to apply for each Design Category and Limit State. The design approach embodied in ASCE 43-05 draws upon, and closely follows, the procedures from the U.S. Department of Energy Standard 1020 (Ref. 4), developed by DOE for high consequence facilities. A very useful document that explains the basis for DOE STD 1020 was published by Kennedy and Short along with the original standard (Ref. 5). The U.S. Nuclear Regulatory Commission's Standard Review Plan NUREG-0800 (Ref. 6) also references ASCE 43-05, although the agency does not endorse all of it for use by licensees and applicants. (See below.)

When a designer begins the design process, the first step is to select the SDC and the Limit State, from which the rest of the requirements and the procedure flow.

<u>*Target Performance Goals*</u>: ASCE 43-05 also provides Target Performance Goals for each SDC. These goals are described in terms of *mean annual probability of exceedance*, and are shown in Table 7. ASCE 43-05 actually only provides the Goals for SDC categories 3, 4, and 5. For categories 1 and 2, the 43-05 standards committee estimated what approximate performance goal is embedded in ASCE 7, but this is not explicitly written down in ASCE 7, and as mentioned categories 1 and 2 are outside of our scope here.

In the ASCE 43-05 procedures, the *Seismic Design Category* is used to establish the design earthquake level motions. The *Limit State* is used to select the appropriate design procedures, analysis methodology, and acceptance criteria.

SSCs in nuclear power plants that perform vital safety functions are assigned in ASCE 43-05 to SDC 5 and Limit State D. The US NRC has adopted this assignment. The Target Performance Goal P_F for SDC 5 is 1×10^{-5} per year. It is this set (SDC 5, LS D, and P_F goal of 1×10^{-5}/year) that will be the focus of our discussion in what follows. Also notice that the approach treats each individual SSC separately – that is, each SSC that falls under the coverage of the standard must individually meet the goal.

TABLE 4
ANSI/ANS 2.26 Seismic Design Category Descriptions for Nuclear SSCs

	Consequences of SSC failure		
	Worker	**Public**	**Environment**
SDC-1	No radiological/toxicological release consequences but failure of SSCs may place facility workers at risk of physical injury	No Radiological/toxicological release consequences	No radiological/toxicological release consequences
SDC-2	Radiological/toxicological exposures to workers will have no permanent health effects, may place more facility workers at risk of physical injury, or may place emergency facility operations at risk	Radiological/toxicological exposures of public areas are small enough to require no public warnings concerning health effects	No radiological or chemical environmental consequences
SDC-3	Radiological/toxicological exposures that may place facility workers long-term health in question	Radiological/toxicological exposures of public areas would not be expected to cause health consequences but may require emergency plans to assure public protection	No long-term environmental consequences are expected, but environmental monitoring may be required for a period of time.
SDC-4	Radiological/ toxicological exposures that may cause long-term health problems and possible loss if life for a worker in proximity of the hazardous material, or place workers in nearby on-site facilities at risk.	Radiological/toxicological exposures that may cause long-term health problems to an individual at the exclusion area boundary for 2 hours	Environmental monitoring required and potential temporary exclusion from selected areas for contamination removal
SDC-5	Radiological/toxicological exposures that may cause loss of life in workers in the facility	Radiological/toxicological exposures that may possible cause loss of life to an individual at the exclusion area boundary for an exposure of 2 hours	Environmental monitoring required and potentially permanent exclusion from selected areas of contamination

TABLE 5
ANSI/ANS 2.26 Description of SSC Limit States

Limit State	Description	Examples
A	An SSC designed to this Limit State may sustain large permanent distortion short of collapse and instability (i.e. uncontrolled deformation under minimal incremental load) but shall still perform its safety function and not impact the safety performance of other SSCs	*Examples*: 1) building structures that must function to permit occupants escape to safety following an earthquake; 2) systems and components designed to be pressure retaining but may perform their safety function even after developing some significant leaks following an earthquake
B	An SSC designed for this Limit State may sustain moderate permanent distortion but shall still perform its safety function. The acceptability of moderate distortion may include consideration of both structural integrity and leak-tightness	*Examples*: 1) building structures that cannot be damaged to the extent that the ability to perform their safety function is lost. Such structures include fire stations, hospitals, or other emergency response structures; 2) systems and components designed to be pressure retaining but may perform their safety function even after developing some minor leaks following an earthquake (i.e. either they do not contain hazardous material, or the leakage rates associated with minor leaks do not exceed the consequence level of the assigned SDC category)
C	An SSC designed to this Limit State may sustain minor permanent distortion but shall still perform its safety function. An SSC that is expected to undergo minimal damage during and following an earthquake such that no post-earthquake repair is necessary may be assigned this Limit State. An SSC in this Limit State may perform its confinement function during and following an earthquake.	*Examples*: 1) glove boxes containing radioactive or hazardous material; 2) confinement barriers for radioactive or hazardous materials; 3) heating ventilation and air-conditioning systems that service equipment or building space containing radioactive or hazardous material; 4) active components that may have to move or change state following the earthquake
D	An SSC designed to this Limit State shall maintain its elastic behavior. An SSC in this Limit State shall perform its safety function during and following an earthquake. Gaseous, particulate and liquid confinement by SSCs is maintained. The component sustains no damage that would reduce its capability to perform its safety function.	*Examples*: 1) containments for large inventories of radioactive to hazardous materials; 2) components that are designed to prevent inadvertent nuclear criticality; 3) SSCs that perform safety functions that may be impaired due to permanent deformation, e.g. valve operations, control rod devices, high-efficiency particulate absorber (HEPA) filter housings, turbine or pump shafts. etc. 4) SSCs that perform safety functions that require the SSC to remain elastic or rigid so that it retains its original strength and stiffness during and following a DBE to satisfy its safety, mission, or operational requirements, e.g. relays, switches valve operators, control rod drives, HEPA filter housings, turbine or pump shafts etc.

TABLE 6
Earthquake Design Provisions and Applicable Standards

Defines analysis methodology and acceptance criteria

Defines earthquake motion input level

SDC	Limit State A	Limit State B	Limit State C	Limit State D
1	ASCE 7	ASCE 7	ASCE 7	NA
2	ASCE 7	ASCE 7	NA	NA
3	ASCE 43-05	ASCE 43-05	ASCE 43-05	ASCE 43-05
4	ASCE 43-05	ASCE 43-05	ASCE 43-05	ASCE 43-05
5	ASCE 43-05	ASCE 43-05	ASCE 43-05	ASCE 43-05

TABLE 7
Target Performance Goals for Seismic Design Categories (SDC) *

SDC	**Target Performance Goal** (Mean annual frequency of unacceptable performance, P_F)
1	$< 1 \times 10^{-3}$
2	$< 4 \times 10^{-4}$
3	$\sim 1 \times 10^{-4}$
4	$\sim 4 \times 10^{-5}$
5	$\sim 1 \times 10^{-5}$

***** ASCE 43-05 only addresses Structural Design Categories 3, 4 and 5. Categories 1 and 2 are covered by existing building codes, and the performance goals for these SDCs were estimated by ASCE Nuclear Working Group members as those approximately achieved by building codes.

5.3 Design guidance in ASCE 43-05

A designer whose task is the design of an SSC to achieve adequate seismic performance needs a specified reference seismic design input to work with. This is usually called a "*design basis earthquake*", although others terms are also used. In laymen's terms, this is the "earthquake motion" that the designer must assure can be withstood. As discussed above, the goal of ASCE 43-05 is to meet a specified target performance goal. While in principle a designer could select his own reference seismic design input and then execute his design using it, and then analyze the design to demonstrate compliance with the target performance goal, that is not the approach used today. Instead, the methodology in ASCE 43-05 is prescriptive. *

To achieve an SSC design that meets the specified target performance goal, a reference seismic hazard defined at annual exceedance frequency H_D is specified. ASCE 43-05 specifies that for SDC 5, H_D is 1 x 10^{-4}/year. This H_D is the starting point. Using this starting point, the designer who follows the rest of the guidance in the standard should produce an SSC that will achieve the target performance goal P_F of 1 x 10^{-5}/year. (This is, however, a target, not an outcome that is absolutely assured in each case.) An overview of how this is accomplished will be described briefly below. The approach is to use adequately conservative deterministic SSC design rules and acceptance criteria, described in the standard.

Note that the choice of H_D as a starting point for achieving P_F is somewhat arbitrary. In principle, one can select H_D within a broad range of hazard exceedances, and still achieve the target P_F. An extreme example might be starting with H_D equal to the target P_F, and using median design properties throughout, or at least embedding much less conservatism in the rest of the design approach. This is extreme because the engineering community generally prefers not to use median design approaches, out of concern that the "tails" of the distributions of the various design features, often not well understood, will compromise the overall conservatism that is sought.

Probability Ratio or Risk Reduction Ratio R_P: [Both of these terms are used in the literature to designate the same ratio quantity, but both of these nomenclatures can be confusing.] The ratio of the hazard annual exceedance frequency H_D to the performance goal annual exceedance frequency P_F is defined as the *probability ratio* R_P, so that

$$R_P = H_D / P_F.$$

This ratio is also sometimes referred to as the *risk reduction ratio* (Ref. 5) since it

* The use of the word "prescriptive" is not exactly true, in the sense that ASCE 43-05, like most other such consensus codes and standards, does provide the user with the option to "do something else", provided that it can be justified. However, this option is seldom exercised in routine practice.

can be thought of as representing an additional risk reduction below the annual risk of exceeding the design basis earthquake. For example, from Table 8, for *Seismic Design Category* 5 SSCs, the Target Performance Goal P_F is 1 x 10^{-5} per year, and the reference *Hazard Exceedance Probability* H_D is 1 x 10^{-4} per year, so the associated probability ratio is 10. This can be viewed as designing the SSC to a 1 x 10^{-4} per year seismic hazard but embedding enough conservatism so that an additional factor of 10 reduction is achieved, leading to at least if not better than a 1 x 10^{-5} per year frequency of "failure", where "failure" means exceeding the Limit State. The probability ratio (or equivalently the risk-reduction ratio) must be achieved by using adequately conservative deterministic SSC design rules. It is emphasized that implicit in the use of the reference hazard H_D is that sufficient conservatism must be embedded in the specified design rules and seismic acceptance criteria to achieve the desired risk reduction ratio and hence the desired target P_F.

TABLE 8
ASCE 43-05 Earthquake Design Provisions

	Seismic Design Category (SDC)		
	3	**4**	**5**
Target Performance Goal Exceedance (P_F) (per year)	1 x 10^{-4}	4 x 10^{-5}	1 x 10^{-5}
Probability Ratio (R_P)	4	10	10
Hazard Exceedance Probability (H_D) (per year)	4 x 10^{-4}	4 x 10^{-4}	1 x 10^{-4}

5.4 Design Basis Earthquake definition – ensuring adequately conservative design rules and acceptance criteria to achieve F_P

In the ASCE 43-05 provisions, it is assumed that the *Design Basis Earthquake* is based upon a site-specific Probabilistic Seismic Hazard Assessment (PSHA) that produces seismic hazard curves and Uniform Hazard Response Spectra (UHRS) associated with several hazard exceedance frequencies, as indicated in Figure 2. Nowadays the most commonly used guidance for carrying out a PSHA is the so-called “SSHAC” guidance, following the methodology established by a DOE-NRC-EPRI expert committee (Ref. 7, 8). Neither the SSHAC process nor the derivation of the UHRS from the H_D following the ASCE 43-05 procedure will be described here.

For Seismic Design Categories 3, 4 and 5, a Uniform Hazard Response Spectrum (UHRS) is specified at the mean Hazard Annual Frequency of Exceedance H_D provided in Table 8. The Design Basis Earthquake (DBE) for evaluation of the SSC is defined in terms of a Design Response Spectrum (DRS) and is given by

$$DRS = DF \times UHRS$$

where the DRS, like the UHRS, is a frequency-dependent spectrum, meaning that the Design Factor DF must also be determined frequency-by-frequency, as shown schematically in Figure 2.

Though not explicitly stated in the main body of ASCE 43-05, the introduction of DF is intended to ensure that the appropriate performance goal is achieved by correcting for the fact that different sites have different slopes of the seismic hazard curve. Developing DF follows a procedure that will not be described here. An abbreviated derivation of the formulation for DF is included in the commentary of ASCE 43-05. A more thorough description of the derivation with mathematical details is included in the development of the basis for DOE standard 1020 (Ref. 5) and in Kennedy (Ref. 9), where a probabilistic analysis is carried out by convolving representative hazard and fragility curves to obtain an analytical expression for the Target Performance Goal P_F. In Kennedy (Ref. 11), an explicit formula for the Design Factor is derived.

5.5 Achieving desired performance – FOSID for NPPs

For nuclear power plants, ASCE 43-05 also selects the targeted design performance conservatively. Specifically, Limit State D (Table 5) is defined so that, “An SSC designed to this Limit State shall maintain its elastic behavior.” This has been interpreted more specifically to be the threshold point for the *“onset of significant inelastic deformation”*. The frequency for reaching this threshold has been given the moniker FOSID (frequency of the onset of significant inelastic deformation).

The US NRC has specifically adopted this procedure for development of site-specific ground motions in Regulatory Guide 1.208 (Ref. 10), where it states:

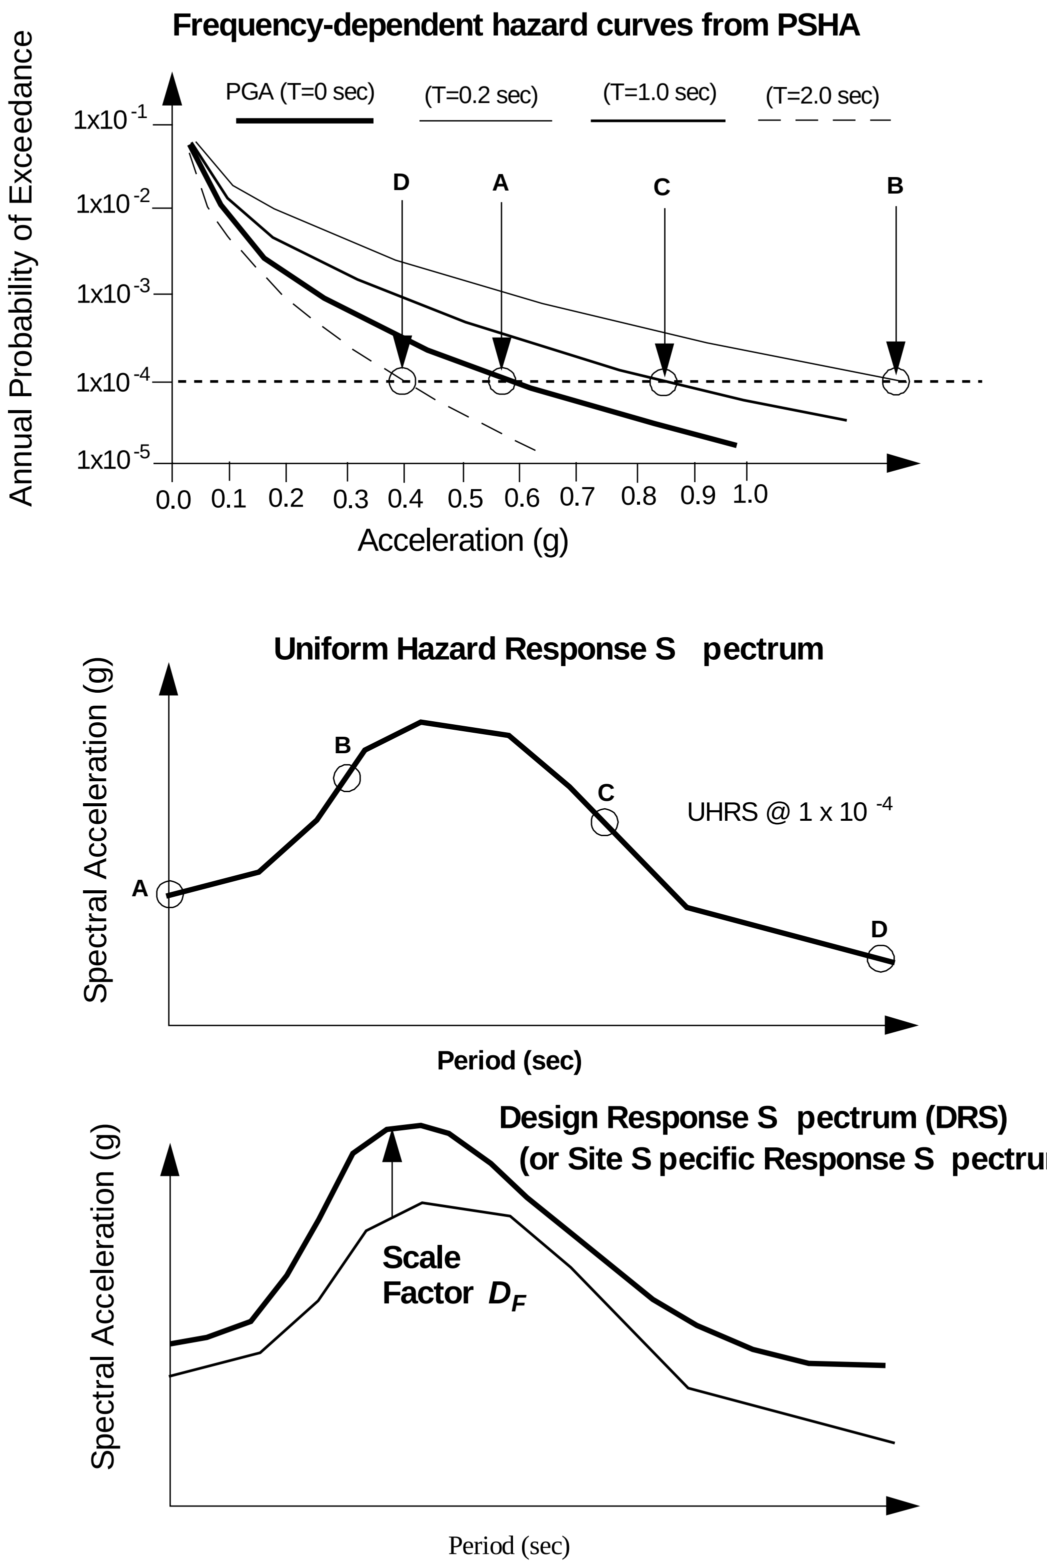

FIGURE 2

Generation of Uniform Hazard Response Spectrum from site-specific hazard curves

"*A performance-based approach is described in Chapters 1 and 2 of ASCE / SEI Standard 43-05, instead of the reference probability approach described in Appendix B to Regulatory Guide 1.165. The performance-based approach employs Target Performance Goal (P_F), Probability Ratio (R_P), and Hazard Exceedance Probability (H_D) criteria to ensure that nuclear power plants can withstand the effects of earthquakes with a desired performance, the desired performance being expressed as the target value of 1×10^{-5} for the mean annual probability of exceedance (frequency) of the onset of significant inelastic deformation (FOSID). This approach targets performance criteria for SSCs which is defined relative to the onset of inelastic deformation; instead of relative to the occurrence of failure of the SSC. The mean annual probability of exceedance associated with this performance target was chosen based on this performance criterion.*"

As noted above, the NRC adoption is equivalent to specifying that the SSCs vital to safety in nuclear power plants are *Structural Design Category 5* SSCs for which only *Limit State D* outcomes are permitted under loads from the Design Basis Earthquake. Thus these SSCs must achieve ruggedness at least sufficient to assure that they achieve an annual probability of exceedance of 1×10^{-5} with respect to FOSID.

It is important to note here that the NRC endorsement of the ASCE 43-05 approach for defining the site-specific earthquake ground motion, although important, does not mean that the NRC has endorsed the rest of the procedures in ASCE 43-05. Only the part of ASCE 43-05 leading to a site-specific design response spectrum has been explicitly endorsed. The NRC, of course, has its own regulations and regulatory guidance, and although much of what the NRC requires or offers as guidance is similar to what is in ASCE 43-05, there are differences in detail, some of which are important although many others are not.

5.6 Achieving desired performance – design provisions

It is beyond our scope here to provide details about how the ASCE 43-05 design specifications and acceptance criteria achieve the desired performance. (As mentioned, much of the technical detail in ASCE 43-05 is based on DOE Standard 1020-2002, Ref. 4.) The technical issues that a designer must be attentive to include, among others, lateral force provisions; story drift/damage control provisions; detailing for ductility provisions; and quality assurance provisions. As Standard 1020 states (Ref. 4), "These provisions are comprised of the following four elements taken together: (1) seismic loading; (2) response evaluation methods, (3) permissible response levels; and (4) ductile detailing provisions."

Details in the code depend on the SSC, of course. Different structural systems are permitted different allowable drift limits, and guidance on allowable inelastic energy absorption factors is intended to assure adequate conservatism and also in part to account for variability. These factors differ for different structural configurations under different levels of axial load. While static analysis is allowed for the lower SDCs and Limit States, for nuclear power plant applications a dynamic analysis is required, for which guidance is provided.

In the evaluation of seismic demand, many different approaches (linear equivalent state analysis, linear dynamic analysis, complex frequency response methods, and nonlinear analysis) are permitted under conditions and constraints specified in the code. Guidance is provided on stiffness, damping values, and other parameters. Taken as a whole, the details, which are beyond our scope here, provide a set of prescriptive design rules that should produce the desired result.

5.7 Performance implications of the ASCE 43-05 design approach

As discussed above, the target performance goal adopted for nuclear power plant applications is a mean annual frequency of 1×10^{-5}/year with respect to the earthquake-caused onset of significant inelastic deformation (FOSID). Crucially, ASCE 43-05 notes that the seismic demand and structural capacity evaluation criteria are aimed at providing sufficient conservatism to achieve two different outcomes for nuclear power plant applications, both of which are stated in *probabilistic terms.* Both of these should be true for any individual SSC that is designed and evaluated according to the ASCE 43-05 framework:

- Less than about a 1% probability of unacceptable performance for the Design Basis Earthquake Ground Motion (with the Design Spectra defined per 43-05) and

- Less than about a 10% probability of unacceptable performance for a ground motion equal to 150% of the Design Basis Earthquake Ground Motion.

The conclusion that these outcomes are achieved if the design rules and acceptance criteria in ASCE 43-05 are followed is based on considerable experience and judgment, backed up by extensive calculations found in Kennedy (Ref. 5, 9).

5.8 Basis for selection of the 1×10^{-5} target performance goal: PRA analysis of current plants

Kennedy notes (Ref. 9) that the basis for selection of the 1×10^{-5} per year target performance goal is related to the calculated seismic contribution to annual core damage frequency in current U.S. nuclear power plants. For the U.S. plants that have completed full Seismic Probabilistic Risk Assessments (SPRAs), which includes about 25 U.S. plants, the average annual seismic contribution to core damage frequency is ~1 x 10^{-5} per year. By selecting the FOSID target (onset of significant inelastic deformation) as the target performance goal at 1 x 10^{-5}/year, there is additional significant conservatism embedded with respect to the potential for core damage, because FOSID corresponds to a significantly lower response level than that required to reach any type of important damage for an individual SSC that might lead to overall reactor core damage.

The Electric Power Research Institute performed a rigorous convolution of fragility and hazard curves for 28 central and eastern U.S. nuclear power plant sites (Ref. 11). Modern PSHA hazard studies were performed for each of these sites in accordance with the EPRI methodology (Ref. 12) and Safe Shutdown Earthquakes were computed for each site in accordance with ASCE 43-05 with a target performance goal of $1x10^{-5}$ with respect to FOSID. The fragility and hazard convolutions indicated that the resulting realized FOSID risks for these plants were in the range 0.54 to 1.07 x 10^{-5}/year.

In addition, Kennedy has also evaluated core damage frequency for a number of representative plants by numerical convolution of the hazard and lognormal fragility curves (Ref. 9). Kennedy’s results indicated core damage frequencies in the range from 0.6 to 4.3 x 10^{-6} per year. Based on the results of these studies, Kennedy has noted that adopting a target goal of FOSID at mean 1 x 10^{-5}/year implies that the Core Damage Frequency will generally be significantly below the target of 1 x 10^{-5}/year and in fact it is expected that the CDF will be in the range of 0.6 x 10^{-6} to 6 x 10^{-6}/year (Ref. 9).

The structural performance implications of utilization of the ASCE standard 43-05, and NUREG-0800 which adopts the ASCE provisions, are summarized in Figure 3.

Applies to individual SSCs at nuclear power plants

Adopts ASCE 43-05 procedure for definition of the Safe Shutdown Design Spectrum

Performance target of 1 x 10^{-5} per year with respect to the onset of significant inelastic behavior

Conservatisms of seismic demands (defined by SSE spectra) and structural capacity criteria (national consensus structural design codes) are sufficiently conservative to achieve:

- Less than about 1% probability of unacceptable performance for the design basis earthquake motions
- Less than about 10% probability of unacceptable performance for ground motion equal to 150% of the design basis earthquake ground motions
-

Realized reactor seismic core damage frequencies expected to be in the range of 0.6 to 6 x 10^{-6} per year

Results in a "risk consistent" design in that plants designed for different sites (where seismic hazard curves have potentially different shapes in the frequency range of interest) will have consistent seismic risks

FIGURE 3

Nuclear power plant seismic performance implications from the utilization of the ASCE 43-05 provisions adopted in NUREG-0800

6.0 Characterization of today's approach: which "options" are represented

6.1 Introduction

In Section 4, a set of "options" was presented for how a "framework" could be constructed, covering the *performance criteria*, the *acceptance criteria* (including a confidence level), *analytical methods*, and *design criteria*. Tables 1, 2, and 3 in Section 4 provided a summary of these options.

Here we will characterize today's "framework", by which we mean the framework that follows ASCE 43-05 and the NRC's corresponding guidance for nuclear power plant safety applications, vis-à-vis these several "options." We will do this based on the overview in Section 5 of today's framework.

A summary of which options are represented by today's framework is shown in Tables 9, 10, and 11, which are identical to Tables 1, 2, and 3 except that the options represented by today's framework are explicitly highlighted.

The next several paragraphs explain how today's framework can be characterized vis-à-vis the several options. The reader is referred to Tables 9, 10, and 11 as a quick visual guide to what follows.

6.2 Performance criterion

As discussed in Section 4, today's framework uses an *annual frequency of unacceptable performance* as its performance criterion. For SSCs vital to the safety of nuclear power plants, this is linked to Seismic Design Category 5, and selected to be a target frequency of 1×10^{-5} per year.

6.3 Confidence level

The approach in today's framework explicitly seeks *mean confidence* that the performance criterion is met. This mean confidence is typically around one standard deviation above the median, and sometimes higher.

6.4 Performance figure-of-merit

The *onset of significant inelastic deformation* as the threshold for "failure" in Limit State D is definitely not a realistic characterization of the earthquake-caused "failure" of the SSCs under consideration. *FOSID is a conservative characterization of failure*, and indeed has been explicitly chosen to be so. As noted in Section 4, the onset of inelastic deformation is conservatively a long way short of the point where a given SSC will fail to perform its safety function.

6.5 End point – a single SSC

The undesired end-point of concern in the seismic "framework" under discussion here is the *seismic-caused failure of a single SSC*. The next paragraph explains why for nuclear power-plant applications this is conservative, and indeed almost always highly so.

The undesired end-point of concern in a nuclear power plant is damage to the core, as analyzed in a PRA. In a seismic PRA (SPRA), the scenarios leading to core damage are all characterized by the initiating earthquake causing damage to one or more SSCs, leading to core damage. Sometimes, depending on the scenario, one or more non-seismic failures or human errors must occur too. In all of the few dozen SPRAs to date that have studied large LWR power plants, no scenario has been found to be important in which core damage is caused by the seismic failure of a single SSC – it always requires more than one, or an SSC failure plus a non-seismic failure or a human error. This of course is largely because the systems design of the plants, with its redundancy, diversity, and defense in depth, explicitly assures that such a "singleton" scenario cannot by itself lead directly to core damage.

Thus, choosing a single SSC as the end point for the "framework" is highly conservative, vis-à-vis the true end-point of safety concern, core damage.

6.6 Realistic or conservative analysis

The analysis that the current "framework" requires to show compliance, as described in Section 4, is a *combination of realistic and conservative analysis.* Certain features of the analysis are explicitly conservative, while others are more nearly realistic.

6.7 Design requirements

The design requirements are *detailed and prescriptive*. As described briefly in Sections 4 and 5, these requirements aim toward an *explicit performance target.* Experience from the seismic PRA literature for large LWRs demonstrates that this design approach is successful, in that the performance target is met by the current fleet of plants.

6.8 Probabilistic performance outcomes

As noted in Section 5, ASCE 43-05 states that if the framework is followed for a specific SSC, sufficient conservatism is embedded that both of the following outcomes will likely be true. Note that both of these outcomes are framed in probabilistic terms:

TABLE 9
Options for the probabilistic performance criterion and for the confidence level

probabilistic performance criterion	confidence level
Annual frequency of unacceptable performance	**explicit high confidence**
	median confidence + extra margin arising from somewhere else
a surrogate	explicit high confidence
	median confidence + extra margin arising from somewhere else

TABLE 10
Options for performance and for analysis

Options: Performance	Options: End-Point	Options: Analysis
Realistic Success v. Failure	full ensemble of accident scenarios	realistic analysis
		conservative analysis
	one accident scenario	realistic analysis
		conservative analysis
	a single SSC	realistic analysis
		conservative analysis
Conservative Success vs. Failure	full ensemble of accident scenarios	realistic analysis
		conservative analysis
	one accident scenario	realistic analysis
		conservative analysis
	a single SSC	**realistic analysis**
		conservative analysis

TABLE 11
Design side options

True Performance-Based Design	**Specified Deterministic Design Rules with an Explicit Performance Criterion**	Specified Deterministic Design Rules without an Explicit Performance Criterion

- Less than about a 1% probability of unacceptable performance for the Design Basis Earthquake Ground Motion (with the Design Spectra defined per 43-05) and

- Less than about a 10% probability of unacceptable performance for a ground motion equal to 150% of the Design Basis Earthquake Ground Motion.

6.9 Summary

As can be seen, the current framework is effective, in that it will produce a nuclear power plant that will meet the performance target with significant margins. These margins are embedded explicitly in some of the aspects of the framework and implicitly in others. The framework is also one that can be used effectively by engineers at the bench who do the design and analysis work, and it can be efficiently and effectively peer-reviewed. All of these features are attractive, and one of the most attractive of them is the explicit numerical performance target.

Shortcomings: There are, however, shortcomings that make the framework less than ideal. None of these get in the way of the effective implementation of today's framework, but some of them, if overcome, could provide a much sounder foundation for the future.

Some possible paths forward are discussed in the next Section.

7.0 Possible paths forward beyond today's "framework"

7.1 Introduction

The current "framework" for seismic SSC design and analysis for nuclear power-plant applications has many strengths, among the most important of which is that it already has demonstrated that it is technically sound and useful in the hands of both the routine designer and the sophisticated analyst. A huge positive attribute is that it is anchored in a technically achievable <u>and analyzable</u> performance goal that is considered "safe enough." Another major attribute is its strong endorsement throughout the community of practitioners, as well as by the regulators.

Therefore, any suggestions for "improvements" must be made with a certain amount of humility, plus a large dose of "look before you leap" skepticism.

Nevertheless, it is easy to see where, at least in principle, some further advances in the framework are available. The problem is that closing the gap between "available" and "achievable" is, in this arena as in many others, both difficult and likely to take a long time.

With these words of caution, we will proceed to discuss elements of several different advances that could address several different limitations that are manifest in the current "framework."

For each of these, one could formulate a research program that could help define how far each of these ideas could realistically be pushed along.

7.2 Limitation: one SSC at a time

A major limitation is that the current framework forces each SSC in the reactor with the requisite safety significance to be designed, analyzed, and approved by the regulator <u>one at a time</u>. But given that there are no earthquake-caused accident sequences that involve only a single SSC, this means that the current approach is surely conservative, and sometimes significantly so.

At least in principle, the remedy for this limitation is to use a full PRA analysis to understand the specific role of each SSC in all of the accident sequences where it participates. With this understanding, it is likely that some SSCs may not require the full robustness (and cost) of meeting every requirement in the framework. For example, for a given SSC it might be that the only important accident sequences in which it participates are characterized by the SSC being in Boolean logical AND with another different SSC which is demonstrably very strong in earthquakes, much stronger than the framework demands. (What this Boolean AND concept means in plain English is that both SSCs would need to fail to cause the accident sequence.) But if the "other" one, say SSC B, is very strong, then there is not necessarily a need for our component, say SSC A, also to be as strong as the framework demands. Perhaps a less stringent requirement for SSC A could still provide overall adequate safety for the reactor.

At least in principle, one could use the PRA to work out how strong SSC A actually needs to be against earthquakes, in terms of a less stringent performance target.

This approach, at least as expressed here, comes down to *using the overall CDF (core damage frequency) for the entire seismic PRA as the figure of merit.* This figure of merit would become a tool to use in modifying the design requirements for some SSCs, based on their individual roles in achieving overall reactor safety. However, in the end *a set of specific design rules, linked to performance targets like those is today's framework, would be necessary for each SSC*. What would be different would be that *the design rules could be tailored to the individual SSC through a tailored performance target*.

A modification to this approach of using the entire PRA could be to examine the role of a given SSC in only those accident sequences to which it contributes, and to study these sequences individually, one by one, not necessarily in the context of the CDF from the entire PRA as a figure-of-merit but in relative isolation – albeit with the context of the entire PRA also kept in mind. A relaxation in design and performance requirements for that SSC would be permitted when it would not materially affect just those sequences. While this may seem like only a modest difference in approach from the broader proposal, it might differ substantially in the following circumstance: Suppose that the full seismic PRA has major uncertainties, meaning that the seismic CDF has major uncertainties, whose origin arises from a technical issue that is not relevant to the few accident sequences to which our SSC contributes. It could be that a robust conclusion can be reached about our SSC in spite of major difficulties with other parts of the PRA. *Allowing the narrower use of the PRA for just a few individual SSCs on a case by case basis could be an intermediate approach*.

7.3 Limitation: FOSID is not failure

A second limitation of the current "framework" is that FOSID, the onset of significant inelastic deformation, is acknowledged not to represent the failure of most SSCs to perform their safety functions. The technical issue that must be wrestled with is that for some SSCs the earthquake "size" (however defined) that actually does compromise the safety performance is not much larger than where FOSID occurs, whereas for other types of SSCs there is quite a large margin. This leads to the obvious suggestion that perhaps advantage can be taken of this fact at least for the latter class of SSCs, while leaving the current framework in place for the former class.

How to differentiate? Clearly this must start with some sort of knowledge, derived in part from data and in part from analysis, that can support careful deliberations on how to use the knowledge to suggest a possible modification, case by case, of today's framework. It should be non-controversial to do the research to gather whatever insights one can on the subject. After that, how to proceed would depend on what is found.

In any event, as a matter of principle, a useful near-term advance would be the adoption by the code committees and/or the regulators of a position that they would be open in practice (not only in principle) to a proposed case-by-case change in the use of FOSID as the definition of "failure", if supported by the facts.

One area where such a change could provide immediate benefit is for the class of SSCs whose behavior in earthquakes is by design to "go inelastic" for a certain "size" earthquake but to retain adequate safety performance even while "going inelastic" unless a much "larger" earthquake were to occur. Again, if adopted by the code committee(s) or the regulators, a philosophy that would recognize that this type of SSC can have an important role in nuclear power plants could "open a door" – open the door to engineering innovations, cost savings, safety improvements, better understanding and/or analyzability, larger safety margins, or some combination. All of these advantages are closed off for this class of SSCs to the extent that the FOSID threshold continues to be the definition of "failure" for all nuclear power plant SSCs, regardless of the true "failure" behavior of the SSC.

7.4 Limitation: reluctance to use a very different H_D, either a higher H_D with a smaller R_P, or vice versa

Today, for SDC 5 and Limit State D, the standard ASCE 43-05 "framework" exhorts the designer to use H_D representing a 1 x 10^{-4}/year hazard along with a "risk reduction factor" R_P of 10, in order for the design to achieve performance at 1 x 10^{-5} per year. The text of ASCE 43-05 quite explicitly allows a designer (applicant) to use another combination of H_D and R_P, but there is a general tendency in the engineering community not to stray very far from the main-line suggested approach set down in any consensus code. This reluctance is especially strong because there is not adequate published guidance on whether quite different combinations of H_D and R_P could lead to a more effective design – perhaps with more margin, or less cost, or more analyzability, or some combination.

It could be very useful if the NRC, or the industry, were to support a research project to explore whether quite different combinations of H_D and R_P could provide useful benefits in some circumstances. It would of course depend on the class of SSC: perhaps benefits could accrue for shear walls but not for large tanks, or vice versa. Nobody has written down enough about this to provide guidance. If in fact the current baseline prescription (H_D = 1 x 10-4, R_P = 10) is best, the research would tell us so. Given the complexity of the parameter space, it would actually be surprising if today's approach turned out to be absolutely the optimum (or close to it) for all classes of SSCs. But today nobody has developed the information.

7.5 The ideal: true (pure) performance-based design

In a true (pure) performance-based design framework, there would be no constraints at all on the designer – no code rules, no code allowables that must be used, no restrictions on how the design could be developed. The only constraint, and it is a crucial one, is that *an analysis, presumably a robust realistic analysis, would be required to demonstrate that the required performance is accomplished*.

This ideal world should probably never be realized in any real-world application, of course! The reasons are several, perhaps the most important of which is that discarding all of the code restrictions and rules also discards a century of design experience developed by the broader engineering community and distilled into the consensus codes and regulations.

Another major reason why this ideal is unapproachable at present is that, except for a very few simple and idealized systems, a robust, realistic, comprehensive analysis with only modest uncertainties and nearly total confidence in its correctness is simply not feasible. The fact is that we are, as an engineering community, dealing with complex real systems for which one cannot at present get close enough to the above ideal to believe we could do away with our current requirements for margin, conservatism, and the like. And crucially, because even "the best of us" can make errors, complete trust in analysis without requiring extra margins is not likely to become a reality any time soon, for any system within the scope being discussed here.

Nevertheless, the advantages of a certain amount of movement in this direction are manifest, and easy to write down (if not to achieve in practice.) The advantages come in at least two categories, in each of which important innovation could be stimulated that is now partially stifled:

- innovation in the form of advanced design concepts, or advanced approaches to achieving better designs with today's concepts.

- innovation in the form of advanced analysis methods, including analysis using simulation and testing working together more effectively.

The stifling of innovation is never a good thing, one would think. True enough. A major effort to develop more steps in the right direction than are now being considered – perhaps only baby steps at first – might let loose a snowballing of innovative ideas not only in design space and analysis space, but also in code-committee/regulatory-philosophy space.

REFERENCES

1. Seismic Design Criteria for Structures, Systems and Components in Nuclear Facilities, American Society of Civil Engineers, Standard ASCE/SEI 43-05, 2005

2. Categorization of Nuclear Facility Structures, Systems, and Components for Seismic Design, American Nuclear Society, Standard ANS-2.26-2004, 2004

3. Minimum Design Loads for Buildings and Other Structures, American Society of Civil Engineers, Standard ASCE/SEI 7-05, 2006

4. Natural Phenomena Hazards Design and Evaluation Criteria for Department of Energy Facilities, Standard DOE-STD-1020-2002, January 2002

5. Basis for Seismic Provisions of DOE-STD-1020, R.P. Kennedy and S.A. Short, Lawrence Livermore National Laboratory Report, UCRL-CR-111478, 1994

6. Standard Review Plan for the Review of Safety Analysis Reports for Nuclear Power Plants, U.S. Nuclear Regulatory Commission, NUREG-0800, March 2007

7. Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts, R.J. Budnitz, G. Apostolakis, D.M. Boore, L.S. Cluff, K.J. Coppersmith, C.A. Cornell, and P.A. Morris, NUREG/CR-6372, Lawrence Livermore National Laboratory, sponsored by the U.S. Nuclear Regulatory Commission, U.S. Department of Energy, and Electric Power Research Institute, 1997

8. Use of Technical Expert Panels: Applications to Probabilistic Seismic Hazard Analysis, R.J. Budnitz, G. Apostolakis, D.M. Boore, L.S. Cluff, K.J. Coppersmith, C.A. Cornell, and P.A. Morris, Journal of Risk Analysis, 18, No. 4, p. 466, 1998

9. Performance-Goal Based (Risk Informed) Approach for Establishing the SSE Site-Specific Response Spectrum for Future Nuclear Power Plants, R.P. Kennedy, Transactions of the 19th International Conference on Structural Mechanics in Reactor Technology (SMiRT), Toronto, Canada, August 2007

10. A Performance-Based Approach to Define the Site-Specific Earthquake Ground Motion, U.S. Nuclear Regulatory Commission Regulatory Guide 1.208, March 2007

11. Performance-Based Seismic Design Spectra, Program on Technology Innovation: Assessment of a Performance-Based Approach for Determining Seismic Ground Motions for New Plant Sites, Electric Power Research Institute, EPRI TR-1-12045, September 2005

12. CEUS (*Central and Eastern U.S.)* Ground Motion Project Final Report, Electric Power Research Institute, EPRI TR-1009684, December, 2004